

**This Page is Inserted by IFW Indexing and Scanning  
Operations and is not part of the Official Record**

**BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** \_\_\_\_\_

**IMAGES ARE BEST AVAILABLE COPY.**

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

L Number	Hits	Search Text	DB	Time stamp
-	1621	redirect\$3 near7 address\$3	USPAT; US-PGPUB	2004/02/24 14:04
-	832	(redirect\$3 near7 address\$3) and @ad<20000908	USPAT; US-PGPUB	2004/08/30 15:49
-	181	((redirect\$3 near7 address\$3) and @ad<20000908) and (server\$1 with client\$1)	USPAT; US-PGPUB	2004/02/26 14:02
-	62	((((redirect\$3 near7 address\$3) and @ad<20000908) and (server\$1 with client\$1)) and balanc\$3	USPAT; US-PGPUB	2004/02/24 14:05
-	28	(((((redirect\$3 near7 address\$3) and @ad<20000908) and (server\$1 with client\$1)) and balanc\$3) and (respon\$6 adj3 time\$1)	USPAT; US-PGPUB	2004/02/24 14:13
-	8	((((((redirect\$3 near7 address\$3) and @ad<20000908) and (server\$1 with client\$1)) and balanc\$3) and (respon\$6 adj3 time\$1)) and (client\$1 adj2 address\$3)	USPAT; US-PGPUB	2004/02/24 14:13
-	1	6665702.pn.	USPAT; US-PGPUB	2004/02/26 12:27
-	203	classless\$3	USPAT; US-PGPUB	2004/02/26 14:32
-	181	((redirect\$3 near7 address\$3) and @ad<20000908) and (server\$1 with client\$1)	USPAT; US-PGPUB	2004/02/26 14:02
-	4	classless\$3 and (((redirect\$3 near7 address\$3) and @ad<20000908) and (server\$1 with client\$1) )	USPAT; US-PGPUB	2004/02/26 14:02
-	832	(redirect\$3 near7 address\$3) and @ad<20000908	USPAT; US-PGPUB	2004/02/26 14:31
-	4	classless\$3 and ((redirect\$3 near7 address\$3) and @ad<20000908 )	USPAT; US-PGPUB	2004/02/26 14:31
-	77	classless\$3 and @ad<20000908	USPAT; US-PGPUB	2004/02/26 14:33
-	887	(redirect\$3 near7 address\$3) and @ad<20000908	USPAT; US-PGPUB	2004/08/30 15:50
-	489	((redirect\$3 near7 address\$3) and @ad<20000908 ) and performance\$1	USPAT; US-PGPUB	2004/08/30 15:50
-	58	(((redirect\$3 near7 address\$3) and @ad<20000908 ) and performance\$1) and cluster\$3	USPAT; US-PGPUB	2004/08/30 15:55

US-PAT-NO: 6415323

DOCUMENT-IDENTIFIER: US 6415323 B1

TITLE: Proximity-based redirection system for robust and scalable service-node location in an internetwork

----- KWIC -----

Application Filing Date - AD (1):

19991209

Brief Summary Text - BSTX (15):

To address the wide variety of problems outlined above, one embodiment of the present invention provides a comprehensive redirection system for content distribution in a virtual overlay broadcast network (OBN). In this system, service nodes are situated at strategic locations throughout the network infrastructure, but unlike previous systems, these service nodes are coordinated across the wide area into a cohesive, coordinated, and managed virtual overlay network. Service node clusters peer with each other across IP tunnels, exchanging routing information, client subscription data, configuration controls, bandwidth provisioning capabilities and so forth. At the same time, the service nodes are capable of processing application-specific requests for content, e.g., they might appear as a Web server or a streaming-media server depending on the nature of the supported service. In short, a service node has a hybrid role: it functions both as a server as well as an application-level content router.

Detailed Description Text - DETX (2):

The comprehensive redirection system of the present invention operates in tandem with service nodes situated at strategic locations throughout the network infrastructure that are coordinated across a wide area into a cohesive, coordinated, and managed virtual overlay network. The overlay network architecture is based on a design philosophy similar to that of the underlying Internet architecture, e.g., it exploits scalable addressing, adaptive routing, hierarchical naming, decentralized administration, and so forth. Because of this, the overlay architecture enjoys the same high degree of robustness, scalability, and manageability evident in the Internet itself. Unlike a physical internetwork, where routers are directly attached to each other over physical links, service nodes in the virtual overlay network communicate with each other using the packet service provided by the underlying IP network. As such, the virtual overlay is highly scalable since large regions of a network (e.g., an entire ISP's backbone) composed of a vast number of individual components (like routers, switches, and links) might require only a small number of service nodes to provide excellent content-distribution performance.

Detailed Description Text - DETX (4):

Fundamentally, service rendezvous entails a system by which it is possible to: (1) publish a single name for a service; (2) replicate the service throughout the network; and (3) have each client that desires the service receive it from the most appropriate server. To scale to millions of clients, the service rendezvous mechanism must efficiently distribute and load-balance client requests to the service nodes spread across the wide area. Moreover, to efficiently utilize network bandwidth, content should flow over the minimum number of network links to reach the requesting client. Both these points argue that clients should be directed to a nearby service node capable of serving the request. If there is no nearby service node capable of servicing the request, the system should be able to redirect the client to a service node elsewhere in the network across the wide area to service the request. Furthermore, it should be possible to cluster service nodes at a particular location and have the clients connect to individual nodes within a cluster based on traffic load conditions. In short, the service rendezvous system should provide a mechanism for server selection and should utilize redirection to effectuate load balancing to achieve the desired result. In cases where a local cluster becomes overloaded, the server selection should compensate to load balance across the wide area.

Detailed Description Text - DETX (16):

The system affords very high availability and robustness when anycast is built on standard adaptive routing protocols and the service elements are clustered for redundancy, thus ensuring that requests are routed only to servers that are properly functioning and advertising their availability; and

Detailed Description Text - DETX (21):

As the Internet and World Wide Web (Web) have grown, ISPs realized that better end-to-end network service performance could be attained by combining two innovative architectural concepts in concert, namely: (1) aggressively peering with a large number of adjacent ISPs at each exchange point; and (2) co-locating data centers containing application services (e.g., Web servers) near these exchange points. The "co-location facility" (colo) at each peering point thus allows application services to be replicated at each peering point so that users almost anywhere in the network enjoy high-speed connections to the nearby service.

Detailed Description Text - DETX (36):

To summarize, the service rendezvous problem is solved in a scalable fashion with two interdependent mechanisms: (1) clients bind to the service infrastructure using anycast addresses and routing; and (2) service nodes bind to the master service site using auxiliary information conveyed explicitly via client URLs or implicitly through a distributed directory like DNS. Excellent scaling performance results by virtue of proximity-based anycast routing and the caching and hierarchy that are built into the DNS.

Detailed Description Text - DETX (41):

In an embodiment of the present invention, a novel scheme called stateful

anycasting is employed. In this approach, the client uses anycast only as part of a redirection service, which by definition, is a short-lived ephemeral transaction. That is, the client contacts an anycast referral node via the anycast service, and the referral node redirects the client to a normally-addressed and routed (unicast) service node. Thus, the likelihood that the redirection process fails because the underlying anycast routes are indeterminate is low. If this does occur, the redirection process can be restarted, either by the client, or depending on context, by the new service node that has been contacted. If the redirection process is designed around a single request and single response, then the client can easily resolve any inconsistencies that arise from anycasting pathologies.

**Detailed Description Text - DETX (53):**

an agent at the termination point for that anycast dialogue redirects the client to a fixed service-node location (i.e., addressed by a standard, non-anycast IP address); and

**Detailed Description Text - DETX (56):**

FIG. 5 shows an embodiment of the present invention that demonstrates how control and service functions are separated within a particular ISP to meet the requirements outlined above. In this embodiment, a service cluster 502 of one or more service nodes (SN) and one or more anycast referral nodes (ARN) are situated on a local-area network segment 504 within a colo 500. The network segment 504 couples to a colo router 506 that in turn, couples to the rest of the ISP and/or the Internet 508.

**Detailed Description Text - DETX (57):**

Under this configuration, a client request 510 from an arbitrary host 512 in the Internet 508 is routed to the nearest ARN 514 using proximity-based anycast routing. The ARN 514 redirects the client (path 516) to a candidate service node 518 (path 520) using the range of techniques described herein. This service model scales to arbitrary client loads because the service nodes are clustered, which allows the system to be incrementally provisioned by increasing the cluster size. In addition, the ARNs themselves can be scaled with local load-balancing devices like layer-4 switches.

**Detailed Description Text - DETX (59):**

There are two key steps to bootstrapping the system: (1) the ARN(s) must discover the existence and addresses of service nodes within the SN cluster; and (2) the ARN(s) must determine which service nodes are available and are not overloaded. One approach is to configure the ARNs with an enumeration of the IP addresses of the service nodes in the service cluster. Alternatively, the system could use a simple resource discovery protocol based on local-area network multicast, where each service node announces its presence on a well-known multicast group and each ARN listens to the group to infer the presence of all service nodes. This latter approach minimizes configuration overhead and thereby avoids the possibility of human configuration errors.

**Detailed Description Text - DETX (62):**

The SN(s) in the service cluster announce their presence and optional information like system load by sending messages to group G.sub.s.

**Detailed Description Text - DETX (66):**

Note that since all the devices in the service cluster are co-located on a single network segment or LAN, the use of IP multicast requires no special configuration of routing elements outside of, or attached to, the LAN.

**Detailed Description Text - DETX (77):**

Having described the local-area architecture of the devices within a single colo installation, a description of a wide-area architecture will be provided which includes how the individual service-node clusters are coordinated and managed across the wide area. There are two main wide-area components for realizing embodiments of the content overlay network included in the present invention, namely:

**Detailed Description Text - DETX (81):**

This disclosure describes how the anycast-based redirection system interfaces with available content delivery systems. A novel framework is used in which service-specific interactions are carried out between the ARN, the SN, the client, and potentially the originating service or content site. For example, the client might initiate a Web request to anycast address A, which is routed to the nearest ARN advertising reachability to address A, which in turn redirects the client to a selected SN with a simple HTTP redirect message, or the Web request may be serviced directly from the ARN.

**Detailed Description Text - DETX (98):**

The ARN selects a candidate service node S from its associated service cluster. The selection decision may be based on load and availability information that is maintained from a local monitoring protocol as described above.

**Detailed Description Text - DETX (106):**

Now, it will be assumed that service node 806 fails. The client 802 notices a disruption in service and reacts by re-invoking the stateful anycast procedure described in the previous section: a service request 830 is sent to the anycast address A and received by the ARN 804, which responds with a redirection message 832, directing the client to a new service node 810. The client can now request the new service feed from the service node 810, as shown at 834. The service node 810 sends a subscription message to the node 808 as shown at 836, and the content again flows to the client as shown at 838. Assuming the client utilizes adequate buffering before presenting the streaming-media signal to the user (as is common practice to counteract network delay variations), this entire process can proceed without any disruption in service. When the client attaches, it can send packet retransmission requests to service node 810 to position the stream appropriately and retransmit only those packets that were lost during the session failover process.

Claims Text - CLTX (19):

6. A method of operating a packet-switched network including addressable routers for routing packet traffic, wherein a packet of data is routed from a source node to a destination node based on address fields of the packet, and wherein the packet-switched network includes a redirector coupled to at least one of the addressable routers and at least one service node, the method comprising:

Claims Text - CLTX (20):

advertising, to an addressable router coupled to the redirector, reachability to an anycast destination address from the redirector, wherein a packet sent to the anycast destination address can be routed to a plurality of destination nodes;

Claims Text - CLTX (21):

accepting a service request from a client at the redirector, wherein the service request is an anycast message to the anycast destination address; and

Claims Text - CLTX (36):

advertising, to an addressable router coupled to the redirector, reachability to an anycast destination address from the redirector, wherein a packet sent to the anycast destination address can be routed to a plurality of destination nodes;